Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-03-16 11:16:35
REFERENCE NO: 183

Author Names & Affiliations

- Paul Zimmerman - University of Michigan

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Chemistry

Title of Submission

quantum chemistry computation

Abstract (maximum ~200 words).

Quantum chemistry and electronic structure theories require computational resources that stretch the abilities of typical compute environments. Demand for high memory, many processor cores, and long walltimes usually makes finding suitable computational platforms difficult.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

Recent work in my research group has pushed the boundaries of electronic structure theory to allow highly parallel computations of exact electronic correlation energies (Zimmerman J. Chem. Phys. 2017). While such simulations are highly parallel, the "unit" of computation that can be parallelized is still a substantial amount of work, demanding ~100GB or more of RAM per MPI process, and walltimes of 12-48 hours. This leaves the code incompatible with novel compute architectures such as GPUs, and many cluster queuing systems simply won't allow jobs that run for 4-5 days, even if they have the proper amounts of memory available. Due to this, I'm forced to run the majority of such simulations on my local hardware, and therefore cannot exploit the intrinsic parallelization to 10,000s of cores due to lack of resources.

While this example is given for my most recent code, I've frequently run into similar situations in using even DFT simulations to optimize an ensemble of reaction pathways (Zimmerman JACS 2016). Here the memory requirements are relatively low, but the absolute independence of the simulations means that even using MPI is overkill -- 1000+ separately queued jobs on 4-8 cores each does the trick, with each job requiring different total walltimes. In this case, queuing systems are configured for a small number of jobs, each on many cores and requiring a relatively fixed walltlime. Again, existing computation resources are poorly suited for quantum chemistry.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

See question 1. Many emerging quantum chemical techniques cannot be easily performed on existing cluster resources.

Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode)."